



US Patent & Trademark Office

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

+partial +backup incremental backup, full backup, partial back



THE ACM DIGITAL LIBRARY



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

**partial backup incremental backup full backup partial backup duplicate backup**

Found 772 of 151,219

Sort results  
by

relevance



[Save results to a Binder](#)

[Try an Advanced Search](#)

[Try this search in The ACM Guide](#)

Display  
results

expanded form



[Search Tips](#)

☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

### 1 [Management of a remote backup copy for disaster recovery](#)

Richard P. King, Nagui Halim, Hector Garcia-Molina, Christos A. Polyzois

May 1991 **ACM Transactions on Database Systems (TODS)**, Volume 16 Issue 2

Full text available: pdf(2.48 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

A remote backup database system tracks the state of a primary system, taking over transaction processing when disaster hits the primary site. The primary and backup sites are physically isolated so that failures at one site are unlikely to propagate to the other. For correctness, the execution schedule at the backup must be equivalent to that at the primary. When the primary and backup sites contain a single processor, it is easy to achieve this property. However, this is harder to do when ...

**Keywords:** database initialization, hot spare, hot standby, remote backup

### 2 [Peer-to-peer infrastructure: Pastiche: making backup cheap and easy](#)

Landon P. Cox, Christopher D. Murray, Brian D. Noble

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue S1

Full text available: pdf(1.65 MB)

Additional Information: [full citation](#), [abstract](#), [references](#)

Backup is cumbersome and expensive. Individual users almost never back up their data, and backup is a significant cost in large organizations. This paper presents *Pastiche*, a simple and inexpensive backup system. Pastiche exploits excess disk capacity to perform peer-to-peer backup with no administrative costs. Each node minimizes storage overhead by selecting peers that share a significant amount of data. It is easy for common installations to find suitable peers, and peers with high ove ...

### 3 [ARIES: a transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging](#)

C. Mohan, Don Haderle, Bruce Lindsay, Hamid Pirahesh, Peter Schwarz

March 1992 **ACM Transactions on Database Systems (TODS)**, Volume 17 Issue 1

Full text available: pdf(5.23 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

DB2TM, IMS, and TandemTM systems. ARIES is applicable not only to database management systems but also to persistent object-oriented languages, recoverable file


systems and transaction-based operating systems. ARIES has been implemented, to varying degrees, in IBM's OS/2TM Extended Edition Database Manager, DB2, Workstation Data Save Facility/VM, Starburst and QuickSilver, and in the University of Wisconsin's EXODUS and Gamma d ...

**Keywords:** buffer management, latching, locking, space management, write-ahead logging

4 A region-based compilation technique for a Java just-in-time compiler

Toshio Suganuma, Toshiaki Yasue, Toshio Nakatani

May 2003 **ACM SIGPLAN Notices , Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation**, Volume 38 Issue 5

Full text available:  pdf(158.62 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Method inlining and data flow analysis are two major optimization components for effective program transformations, however they often suffer from the existence of rarely or never executed code contained in the target method. One major problem lies in the assumption that the compilation unit is partitioned at method boundaries. This paper describes the design and implementation of a region-based compilation technique in our dynamic compilation system, in which the compiled regions are selected a ...

**Keywords:** dynamic compilers, on-stack replacement, partial inlining, region-based compilation

5 Recovery Techniques for Database Systems

Joost S. M. Verhofstad


June 1978 **ACM Computing Surveys (CSUR)**, Volume 10 Issue 2

Full text available:  pdf(2.32 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

6 Highly available systems for database applications

Won Kim

March 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 1


Full text available:  pdf(2.43 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

As users entrust more and more of their applications to computer systems, the need for systems that are continuously operational (24 hours per day) has become even greater. This paper presents a survey and analysis of representative architectures and techniques that have been developed for constructing highly available systems for database applications. It then proposes a design of a distributed software subsystem that can serve as a unified framework for constructing database applica ...

7 A principle for resilient sharing of distributed resources

Peter A. Alsberg, John D. Day

October 1976 **Proceedings of the 2nd international conference on Software engineering**

Full text available:  pdf(749.04 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

A technique is described which permits distributed resources to be shared (services to be offered) in a resilient manner. The essence of the technique is to a priori declare one of the server hosts primary and the others backups. Any of the servers can perform the primary duties. Thus the role of primary can migrate around the set of servers. The concept of n-


host resiliency is introduced and the error detection and recovery schemes for two-host resiliency are presented. The single primary, ...

**Keywords:** Distributed computer systems, Distributed control, Resilient protocols, Resilient resource sharing, Resource sharing

8 High speed on-line backup when using logical log operations

David B. Lomet

May 2000 **ACM SIGMOD Record , Proceedings of the 2000 ACM SIGMOD international conference on Management of data**, Volume 29 Issue 2


Full text available:  [pdf\(220.69 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Media recovery protects a database from failures of the stable medium by maintaining an extra copy of the database, called the backup, and a media recovery log. When a failure occurs, the database is "restored" from the backup, and the media recovery log is used to roll forward the database to the desired time, usually the current time. Backup must be both fast and "on-line", i.e. concurrent with on-going update activity. Conventional online backup sequentially copies ...

9 Programming languages for distributed computing systems

Henri E. Bal, Jennifer G. Steiner, Andrew S. Tanenbaum

September 1989 **ACM Computing Surveys (CSUR)**, Volume 21 Issue 3


Full text available:  [pdf\(6.50 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

When distributed systems first appeared, they were programmed in traditional sequential languages, usually with the addition of a few library procedures for sending and receiving messages. As distributed applications became more commonplace and more sophisticated, this ad hoc approach became less satisfactory. Researchers all over the world began designing new programming languages specifically for implementing distributed applications. These languages and their history, their underlying pr ...

10 File servers for network-based distributed systems

Liba Svobodova

December 1984 **ACM Computing Surveys (CSUR)**, Volume 16 Issue 4

Full text available:  [pdf\(4.23 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#), [review](#)

11 An efficient primary-segmented backup scheme for dependable real-time communication in multihop networks

Krishna Phani Gummadi, Madhavarapu Jnana Pradeep, C. Siva Ram Murthy

February 2003 **IEEE/ACM Transactions on Networking (TON)**, Volume 11 Issue 1

Full text available:  [pdf\(606.18 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Several distributed real-time applications (e.g., medical imaging, air traffic control, and video conferencing) demand hard guarantees on the message delivery latency and the recovery delay from component failures. As these demands cannot be met in traditional datagram services, special schemes have been proposed to provide timely recovery for real-time communications in multihop networks. These schemes reserve additional network resources (spare resources) *a priori* along a backup channel ...

**Keywords:** backup channel, backup multiplexing, dependable connection, multihop network, primary channel, quality-of-service (QoS), real-time communication, resource

reservation protocol (RSVP), segmented backup

12 Industrial sessions: beyond relational tables: Coordinating backup/recovery and data consistency between database and file systems

Suparna Bhattacharya, C. Mohan, Karen W. Brannon, Inderpal Narang, Hui-I Hsiao, Mahadevan Subramanian

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available:  pdf(1.44 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


Managing a combined store consisting of database data and file data in a robust and consistent manner is a challenge for database systems and content management systems. In such a hybrid system, images, videos, engineering drawings, etc. are stored as files on a file server while meta-data referencing/indexing such files is created and stored in a relational database to take advantage of efficient search. In this paper we describe solutions for two potentially problematic aspects of such a data ...

**Keywords:** DB2, content management, database backup, database recovery, datalinks

13 Comparison of Backup Products

Charles Curley


October 2000 **Linux Journal**

Full text available:  html(24.81 KB) Additional Information: [full citation](#), [references](#), [index terms](#)

14 Effectiveness of cross-platform optimizations for a java just-in-time compiler

Kazuaki Ishizaki, Mikio Takeuchi, Kiyokuni Kawachiya, Toshio Suganuma, Osamu Gohda, Tatsushi Inagaki, Akira Koseki, Kazunori Ogata, Motohiro Kawahito, Toshiaki Yasue, Takeshi Ogasawara, Tamiya Onodera, Hideaki Komatsu, Toshio Nakatani

October 2003 **ACM SIGPLAN Notices , Proceedings of the 18th annual ACM SIGPLAN conference on Object-oriented programing, systems, languages, and applications**, Volume 38 Issue 11

Full text available:  pdf(405.65 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


This paper describes the system overview of our Java Just-In-Time (JIT) compiler, which is the basis for the latest production version of IBM Java JIT compiler that supports a diversity of processor architectures including both 32-bit and 64-bit modes, CISC, RISC, and VLIW architectures. In particular, we focus on the design and evaluation of the cross-platform optimizations that are common across different architectures. We studied the effectiveness of each optimization by selectively disabling ...

**Keywords:** Java, just-in-time compiler, optimization

15 Collaborative backup for dependable mobile applications

Marc-Olivier Killijian, David Powell, Michel Banâtre, Paul Couderc, Yves Roudier

October 2004 **Proceedings of the 2nd workshop on Middleware for pervasive and ad-hoc computing**


Full text available:  pdf(265.85 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We describe the work we are conducting on new middleware services for dependable and secure mobile systems. This work is based on approaches à la peer-to-peer in order to circumvent the problems introduced by the lack of infrastructure in self-organizing networks

of mobile nodes, such as MANETs. The mechanisms we propose are based on collaboration between peer mobile devices to provide middleware services such as trust management and critical data storage. This short paper gives a brief ...


**Keywords:** collaboration, data back-up, mobile applications

- 16 Bayeux: an architecture for scalable and fault-tolerant wide-area data dissemination  
Shelley Q. Zhuang, Ben Y. Zhao, Anthony D. Joseph, Randy H. Katz, John D. Kubiatowicz  
January 2001 **Proceedings of the 11th international workshop on Network and operating systems support for digital audio and video**

Full text available:  [pdf\(272.26 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The demand for streaming multimedia applications is growing at an incredible rate. In this paper, we propose Bayeux, an efficient application-level multicast system that scales to arbitrarily large receiver groups while tolerating failures in routers and network links. Bayeux also includes specific mechanisms for load-balancing across replicate root nodes and more efficient bandwidth consumption. Our simulation results indicate that Bayeux maintains these properties while keeping transmission ...


- 17 Automatic tiling of iterative stencil loops  
Zhiyuan Li, Yonghong Song  
November 2004 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 26 Issue 6

Full text available:  [pdf\(947.69 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Iterative stencil loops are used in scientific programs to implement relaxation methods for numerical simulation and signal processing. Such loops iteratively modify the same array elements over different time steps, which presents opportunities for the compiler to improve the temporal data locality through loop tiling. This article presents a compiler framework for automatic tiling of iterative stencil loops, with the objective of improving the cache performance. The article first presents a ...


**Keywords:** Caches, loop transformations, optimizing compilers

- 18 Stable Internet routing without global coordination  
Lixin Gao, Jennifer Rexford  
June 2000 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2000 ACM SIGMETRICS international conference on Measurement and modeling of computer systems**, Volume 28 Issue 1

Full text available:  [pdf\(1.07 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Border Gateway Protocol (BGP) allows an autonomous system (AS) to apply diverse local policies for selecting routes and propagating reachability information to other domains. However, BGP permits ASes to have conflicting policies that can lead to routing instability. This paper proposes a set of guidelines for an AS to follow in setting its routing policies, without requiring coordination with other ASes. Our approach exploits the Internet's hierarchical structure and the commercial relationship ...

- 19 The TickerTAIP parallel RAID architecture  
Pei Cao, Swee Boon Lin, Shivakumar Venkataraman, John Wilkes  
August 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 3

Full text available:  [pdf\(2.04 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


Traditional disk arrays have a centralized architecture, with a single controller through which all requests flow. Such a controller is a single point of failure, and its performance limits the maximum number of disks to which the array can scale. We describe TickerTAIP, a parallel architecture for disk arrays that distributes the controller functions across several loosely coupled processors. The result is better scalability, fault tolerance, and flexibility. This article presents ...

**Keywords:** RAID disk array, decentralized parity calculation, disk scheduling, distributed controller, fault tolerance, parallel controller, performance simulation

## 20 Comparison of access methods for time-evolving data

Betty Salzberg, Vassilis J. Tsotras

June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2

Full text available:  pdf(529.53 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper compares different indexing techniques proposed for supporting efficient access to temporal data. The comparison is based on a collection of important performance criteria, including the space consumed, update processing, and query time for representative queries. The comparison is based on worst-case analysis, hence no assumptions on data distribution or query frequencies are made. When a number of methods have the same asymptotic worst-case behavior, features in the methods tha ...

**Keywords:** I/O performance, access methods, structures, temporal databases

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	11	(partial adj2 backup\$3) and (backup adj2 server)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:21
L2	441	increment\$3 adj2 backup\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:21
L3	6	L1 and L2	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:21
L4	24292	"711"/.cccls.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:22
L5	6	L1 and L2	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:22
L6	2	L5 and L4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:22
L7	2197	Mcbride.in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:22
L8	52840	Miller.in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:22
L9	53835	Pierre.in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:22

L10	6	(partial near2 backup\$3) and (duplicat\$3 near2 backup\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:23
L11	152	duplicat\$3 adj backup\$3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:24
L12	1	(partial adj2 backup\$3) and (duplicat\$3 adj backup\$3)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:24
L17	1210	backup adj2 server	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:24
L18	11	(partial adj2 backup\$3) and (backup adj2 server)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:25
L19	3	10 and 17	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/03/27 15:25